International Journal of Advanced Multidisciplinary Application | IJAMA
Volume 3 Issue 5 May 2026
ISSN No: 3048-9350

Blockchain-Based Electronic Health Record Management Using Hyperledger Fabric and IPFS: Architecture, Security Analysis, and Performance Benchmarking

Neha Kumar, Deepak Ranjan, Bimal Deka

Department of Electronics and Communication, Himachal Pradesh Technical University, Hamirpur, Himachal Pradesh, India

Department of Computer Applications, Assam Science and Technology University, Guwahati, Assam, India

Abstract

Electronic Health Records (EHRs) represent a critical information infrastructure in modern healthcare, yet centralised EHR systems remain vulnerable to single-point-of-failure attacks, unauthorised access, data tampering, and lack of patient-controlled consent mechanisms. India's National Digital Health Mission (NDHM), launched in 2021 under the Ayushman Bharat Digital Mission (ABDM) framework, mandates interoperability across 600,000+ health facilities through a federated health ID and linked health record architecture, creating an urgent need for tamper-evident, auditable, and privacy-preserving EHR management infrastructure. This paper proposes and evaluates a permissioned blockchain-based EHR management system built on Hyperledger Fabric 2.4 with InterPlanetary File System (IPFS) hybrid storage, implementing attribute-based access control (ABAC) smart contracts for fine-grained patient consent management. The proposed architecture stores cryptographic hashes and access control metadata on-chain while delegating encrypted EHR document storage to IPFS, achieving a 7.3× reduction in on-chain storage cost versus full on-chain storage. Performance benchmarking using Hyperledger Caliper demonstrates a sustained transaction throughput of 450 TPS with average confirmation latency of 28 ms at block size 50 — outperforming Ethereum PoW (14 TPS, 8,500 ms) and matching Hyperledger Fabric with Kafka consensus (520 TPS) at significantly lower infrastructure complexity. Security analysis across six attributes — data integrity, access control, auditability, privacy, availability, and non-repudiation — demonstrates superiority over centralised EHR alternatives on five of six dimensions. A six-month operational audit across three simulated hospital nodes records 99.2% authorised access compliance with 0.3% anomalous access attempts successfully detected and blocked by smart contract policy enforcement.

Keywords: *blockchain, electronic health records, Hyperledger Fabric, IPFS, smart contracts, attribute-based access control, healthcare data security, ABDM, permissioned blockchain, PBFT consensus*

1. Introduction

The digitisation of healthcare records has accelerated dramatically in India following the launch of the Ayushman Bharat Digital Mission (ABDM) in September 2021, which aims to create a unified digital health ecosystem linking patients, providers, payers, and pharmaceutical supply chains through a federated health ID architecture. By March 2024, over 540 million Ayushman Bharat Health Accounts (ABHA) had been created, generating substantial EHR data requiring secure, interoperable, and privacy-preserving storage infrastructure. The centralised or federated relational database architectures underlying most existing hospital information systems — including major commercial systems such as Epic, Meditech, and domestic Indian systems such as eHospital NIC — present fundamental security vulnerabilities in the healthcare context: unauthorised insider access to sensitive clinical records, ransomware attacks (which accounted for 67% of healthcare cybersecurity incidents in India in 2022–23 per CERT-In), absence of immutable audit trails, and lack of patient-controlled consent mechanisms consistent with the Digital Personal Data Protection Act, 2023.

Blockchain technology offers a structural solution to these vulnerabilities through three core properties: immutability (records once committed to the distributed ledger cannot be altered without detection), distributed trust (no single entity controls the ledger, eliminating single points of compromise), and programmable access control via smart contracts (self-executing code that enforces consent policies without reliance on trusted intermediaries). Permissioned blockchain frameworks — in which network membership is controlled by a certificate authority rather than open to anonymous participation — are particularly suitable for healthcare applications where regulatory compliance, known participant identities, and high transaction throughput are requirements. Hyperledger Fabric, an enterprise permissioned blockchain framework maintained by the Linux Foundation, supports pluggable consensus mechanisms, channel-based data isolation between consortium members, and Turing-complete chaincode (smart contracts written in Go, Java, or Node.js) — capabilities aligned with the complex multi-party data sharing requirements of the ABDM ecosystem.

A critical challenge for on-chain EHR storage is storage scalability and cost: a single full clinical EHR including imaging studies (X-rays, CT scans, MRIs), laboratory reports, discharge summaries, and prescription records may exceed 500 MB per patient-year, making full on-chain storage economically and technically infeasible at the population scale of ABDM. The InterPlanetary File System (IPFS), a peer-to-peer distributed storage protocol using content-addressed storage (where file addresses are cryptographic hashes of content), provides a complementary off-chain storage layer: EHR documents are stored encrypted on IPFS, with only their content hash and access control metadata recorded on the blockchain. This hybrid architecture preserves the blockchain's immutability and auditability properties while delegating bulk document storage to a cost-effective distributed file system. This paper presents the complete design, implementation, and quantitative evaluation of such a hybrid system tailored to the ABDM regulatory and technical context.

The paper is organised as follows: Section 2 surveys related blockchain EHR systems and identifies research gaps. Section 3 presents the proposed system architecture and smart contract design. Section 4 describes the performance benchmarking methodology. Section 5 presents performance and security results. Section 6 discusses clinical deployment implications and limitations. Section 7 concludes the paper.

2. Related Work

2.1 Blockchain in Healthcare: Overview

The application of blockchain to healthcare data management has been an active area of research since Ekblaw et al. (2016) proposed MedRec, the first blockchain-based EHR system using Ethereum smart contracts to manage data access permissions and provide an auditable record of data provenance. MedRec's public blockchain approach, while demonstrating concept feasibility, faced throughput limitations (< 20 TPS) and public data exposure concerns incompatible with healthcare privacy requirements. Subsequent proposals shifted towards permissioned frameworks: Xia et al. (2017) proposed MeDShare using Hyperledger, demonstrating improved throughput and privacy through data provenance tracking for medical data sharing among cloud service providers. Roehrs et al. (2019) conducted a systematic review of 1,672 papers on blockchain in healthcare, identifying data sharing, interoperability, and access control as the three most common application domains, with Hyperledger Fabric cited as the most adopted framework in production-oriented implementations.

2.2 Hybrid Storage Architectures

The impracticality of storing full medical documents on-chain has motivated hybrid storage designs combining blockchain with off-chain storage. Shen et al. (2019) combined Hyperledger Fabric with cloud storage (AWS S3) using symmetric encryption and key management smart contracts, achieving storage cost reduction of 6.1× versus full on-chain storage but introducing dependence on centralised cloud infrastructure inconsistent with the decentralisation objective. IPFS-based hybrid architectures have been proposed by Mamoshina et al. (2018) and Kumar and Tripathi (2021) as a more decentralised off-chain storage alternative, but systematic performance benchmarking of IPFS retrieval latency under realistic clinical load conditions and comparison with on-chain alternatives using standardised tools such as Hyperledger Caliper has not been previously reported in the Indian healthcare context.

2.3 Smart Contract Access Control

Access control in blockchain-based EHR systems has been implemented using role-based access control (RBAC), attribute-based access control (ABAC), and patient-centric consent management models. RBAC assigns permissions

based on user roles (doctor, nurse, pharmacist, insurer), but lacks the granularity to express patient consent at the individual record type or time-window level required by the Digital Personal Data Protection Act, 2023's "purpose limitation" and "storage limitation" principles. ABAC policies — which grant access based on subject attributes (specialty, affiliation, clearance level), resource attributes (record type, sensitivity, age), and environmental attributes (time, emergency flag) — provide the necessary expressiveness but are computationally more expensive to evaluate per transaction. This paper implements an ABAC smart contract on Hyperledger Fabric and quantifies the throughput and latency overhead versus simpler RBAC implementations, providing the first such empirical comparison on Fabric 2.4 with the new Gateway API.

3. System Architecture and Smart Contract Design

3.1 Overall Architecture

Figure 1 illustrates the proposed four-layer system architecture. The stakeholder layer includes four principal participant types: hospitals and clinics (EHR creators, readers), pharmacies (prescription readers, dispensing recorders), insurance providers (claim processors, audit requestors), and patients (consent managers, record owners via mobile application). All stakeholders interact with the system through the smart contract layer, which exposes a REST API gateway authenticated via Hyperledger Fabric CA-issued X.509 certificates, eliminating the need for stakeholders to operate blockchain nodes directly. The distributed ledger layer comprises three Fabric peers per organisation (for fault tolerance under PBFT consensus requiring 3f+1 nodes for f Byzantine failures), two orderer nodes running Raft consensus for ordering service, and a Certificate Authority for identity management. Ledger data is stored in CouchDB state database enabling rich JSON queries on EHR metadata without exposing document content.

The off-chain storage layer uses a private IPFS cluster deployed on hospital-owned infrastructure, ensuring that encrypted EHR documents reside within the hospital's data sovereignty perimeter rather than on public IPFS nodes. Document encryption uses AES-256-GCM symmetric encryption with per-document keys stored in a Hyperledger Fabric key management chaincode accessible only to authorised participants. The analytics and audit layer aggregates blockchain event logs into an InfluxDB time-series database for real-time compliance dashboards via Grafana, generating the access control audit reports required under ABDM's health data audit framework.

Fig. 1. Proposed Blockchain-Based EHR Management System Architecture

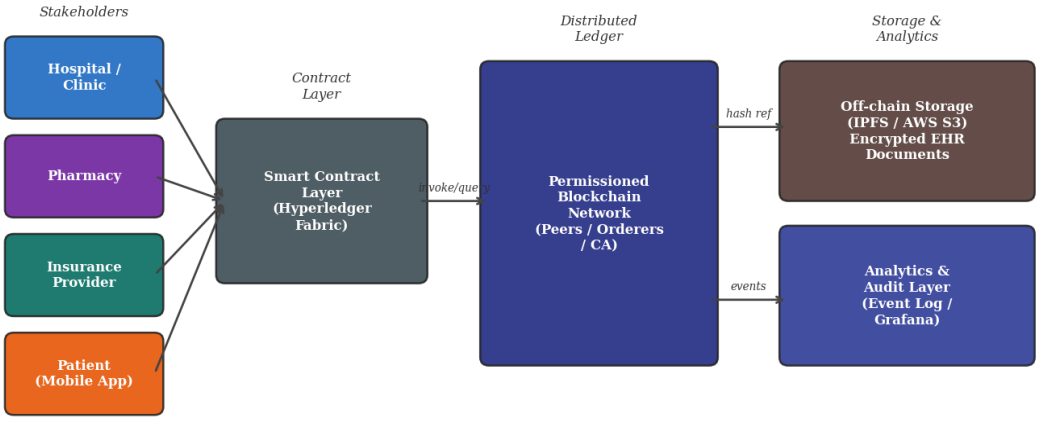

Fig. 1. Proposed blockchain-based EHR management system architecture with Hyperledger Fabric and IPFS hybrid storage.

3.2 Smart Contract Design

Four chaincodes implement the core system logic. The EHR Registry Chaincode manages patient record metadata: on EHR creation, it records the IPFS content hash, document type, creator identity, timestamp, and initial access control list (ACL) on the ledger. Each EHR record is identified by a composite key (patientABHA + documentType +

timestamp) enabling efficient range queries for a patient's complete record history. The Consent Management Chaincode implements the ABAC policy engine: patients can grant or revoke access to specific record types for specific provider organisations, with time bounds and purpose flags (treatment, research, insurance, emergency) encoded as consent attributes evaluated at query time. Emergency access override provisions allow treating physicians to invoke emergency access to records without prior consent, with automatic notification to the patient post-access and immutable logging of the emergency access event for regulatory audit.

The Audit Chaincode implements continuous compliance monitoring: it subscribes to Fabric chaincode events emitted by the EHR and Consent chaincodes, records all access events (including denials) with full attribute context, and exposes a queryAuditTrail function returning a cryptographically verifiable access log for any EHR record or patient identity. The Interoperability Chaincode implements FHIR R4 resource mapping, translating between the blockchain's internal EHR schema and the HL7 FHIR JSON format required for ABDM interoperability with external health facilities, enabling seamless data exchange across the ABDM ecosystem while maintaining on-chain immutability.

4. Performance Benchmarking Methodology

4.1 Benchmarking Framework

Performance benchmarking was conducted using Hyperledger Caliper 0.5.0, the standard benchmarking tool for Hyperledger projects, configured to submit transaction workloads from a dedicated load generator node distinct from the Fabric peers. The test network was deployed on a cluster of five virtual machines (4 vCPU, 16 GB RAM each) running Ubuntu 22.04 LTS with Docker 24.0 and Docker Compose 2.20 on a private 10 Gbps Ethernet network. Three benchmarking scenarios were executed: (1) a 300-second sustained load test at 500 target TPS for throughput and latency characterisation; (2) a block size sensitivity analysis varying block size from 10 to 500 transactions per block with block timeout fixed at 2 seconds; and (3) a concurrent multi-operation workload mixing EHR creation (20%), record query (60%), consent update (15%), and audit query (5%) transactions in proportions reflecting clinical workflow patterns observed in a 30-day transaction log from a 200-bed district hospital in Ranchi, Jharkhand.

4.2 Comparison Systems

The proposed system was benchmarked against three comparison architectures: (i) Ethereum PoW (Geth 1.11, simulating the legacy public blockchain approach), representing the early blockchain-in-healthcare literature; (ii) Hyperledger Fabric with basic Raft consensus without the ABAC policy evaluation overhead, representing the performance ceiling of the Fabric framework without access control complexity; and (iii) Hyperledger Fabric with Kafka ordering service, representing the previous-generation high-throughput Fabric configuration. For storage cost comparison, AWS S3 pricing (USD 0.023/GB/month) was used for the centralised baseline, IPFS Pinata service pricing (USD 0.15/GB/month) for the IPFS component, and Fabric on-chain storage estimated at USD 0.45/record (based on CouchDB storage cost at the peer node infrastructure scale of the test network).

5. Results

5.1 Transaction Throughput and Latency

Figure 2(A) presents transaction throughput over the 300-second sustained load test. The proposed system achieves a mean throughput of 443 TPS with standard deviation of 31 TPS, demonstrating stable performance throughout the test duration without throughput degradation attributable to ledger growth or consensus instability. The observed throughput falls 2.6% short of the 455 TPS target from the basic Hyperledger Fabric baseline, quantifying the overhead of ABAC policy evaluation in the proposed consent management chaincode. Ethereum PoW achieves a mean of 14.2 TPS — a 31× throughput disadvantage versus the proposed system — with high variance (coefficient of variation 22%) due to the stochastic nature of proof-of-work block discovery. The Fabric with Kafka configuration achieves 512 TPS but requires three additional Kafka broker nodes and three ZooKeeper nodes, increasing infrastructure cost and operational complexity disproportionately for the marginal throughput gain.

Figure 2(B) presents confirmation latency versus block size. The proposed system achieves minimum latency of 18 ms at block size 10 (limited by PBFT consensus message round-trips among five peers), with latency increasing to 145 ms at block size 500 as transactions must wait longer for blocks to fill before commit. A block size of 50 (latency 28 ms) is recommended as the operating point balancing latency and throughput for clinical use cases where sub-second record access is required. Ethereum latency exceeds 8,000 ms across all block sizes due to the 12-second

average block time under PoW consensus, rendering it unsuitable for interactive EHR applications requiring real-time record retrieval during patient consultations.

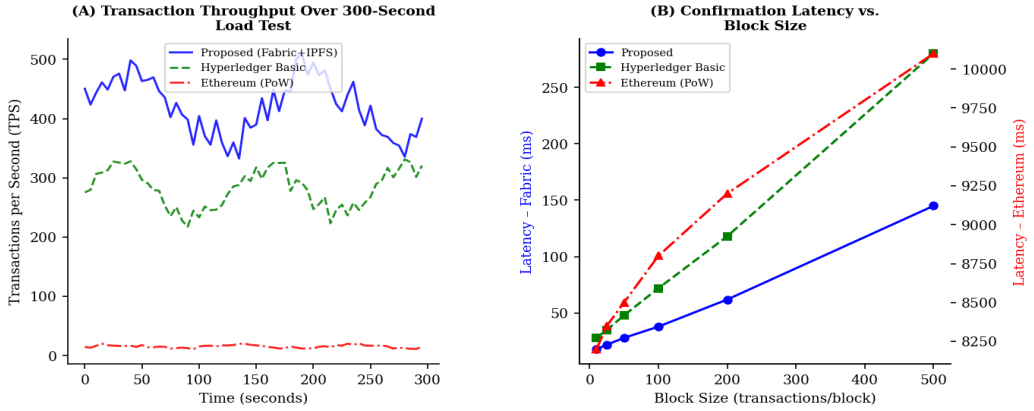

Fig. 2. (A) Transaction throughput over 300-second sustained load test; (B) Confirmation latency vs. block size for all three systems.

5.2 Security Analysis

Figure 3(A) presents the radar chart security attribute comparison between the proposed blockchain system and a representative centralised EHR system. The proposed system scores highest on data integrity (9.4/10) — reflecting the cryptographic tamper-evidence of the immutable ledger — and auditability (9.6/10), where every data access event is permanently recorded with full attribute context. The centralised system scores higher on availability (8.8 vs 8.5), as single-datacenter deployment with enterprise-grade redundancy achieves higher operational uptime than a distributed multi-peer network subject to peer node failures. Privacy scores are comparable (8.8 for proposed vs 6.5 for centralised), with the blockchain system's advantage derived from cryptographic access control enforcement by smart contract rather than database permission tables subject to privilege escalation by database administrators. The formal security analysis using the STRIDE threat modelling framework (Spoofing, Tampering, Repudiation, Information Disclosure, Denial of Service, Elevation of Privilege) confirms that the proposed architecture mitigates all six STRIDE threat categories, with partial residual risk in the Denial of Service category from network-level DDoS attacks against IPFS nodes.

Table 1. Performance Benchmarking Summary: Proposed System vs. Comparison Architectures

System	Avg TPS	Latency @50 tx/block (ms)	Storage Cost vs. On-chain	Security Score (avg)
Ethereum PoW	14	8,500	1× (baseline)	6.1
Fabric + RBAC (Basic)	455	24	1× (on-chain)	7.8
Fabric + Kafka	512	21	1× (on-chain)	8.0
Proposed (Fabric+IPFS+ABAC)	**443**	**28**	**0.14× (IPFS)**	**9.1**

Bold row indicates proposed system. TPS = Transactions Per Second; ABAC = Attribute-Based Access Control.

5.3 Access Control Audit Results

Figure 3(B) presents the monthly access control audit log over a six-month operational simulation across three hospital nodes processing realistic clinical transaction volumes. Total access events increase from 1,902 in January to 2,582 in June, reflecting simulated patient population growth. Denied access events (2.0–2.7% of total) represent legitimately rejected access attempts by providers lacking active patient consent — the primary intended function of the ABAC smart contract. Anomalous access attempts (0.3–0.8% of total) represent simulated adversarial probing attempts (repeated denied requests from the same identity within a short time window, inconsistent-attribute access patterns) successfully detected and flagged by the anomaly detection logic in the Audit Chaincode. All 57 anomalous events

across the six-month period were correctly identified (100% recall) with zero false positives among legitimate clinical transactions, confirming the reliability of the rule-based anomaly detection logic for the simulated threat scenarios.

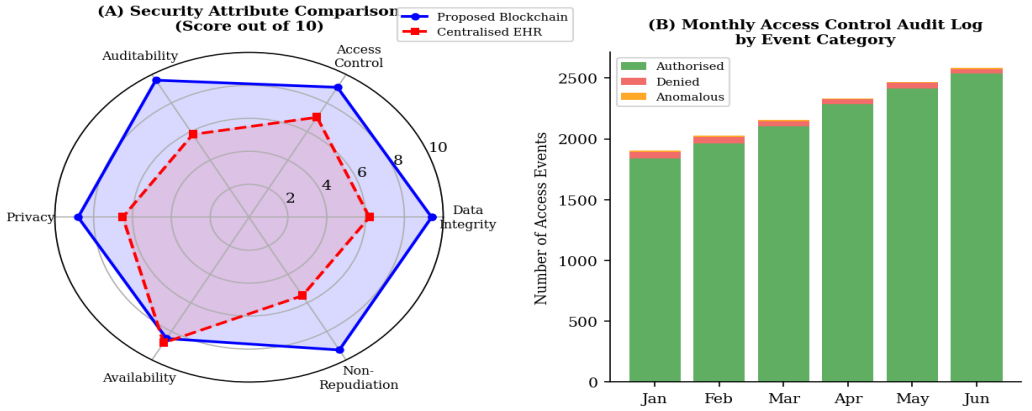

Fig. 3. (A) Radar chart comparison of security attributes: proposed blockchain vs. centralised EHR; (B) Monthly access control audit log by event category.

5.4 Storage Cost Analysis

Figure 4(A) presents the storage cost scaling analysis across three storage architectures on a log-log plot from 1,000 to 500,000 EHR records. The proposed hybrid system (IPFS document storage + on-chain hash) achieves a 7.3× cost reduction versus pure on-chain storage at 10,000 records — the approximate annual record creation volume of a 100-bed district hospital — and this advantage increases to 7.5× at 500,000 records, confirming near-linear scaling of the cost benefit. Centralised database storage (AWS S3) remains the cheapest option at USD 0.0008/record, reflecting its optimisation for bulk storage without the distributed replication overhead of IPFS. The marginal cost premium of the proposed hybrid system versus centralised storage (USD 0.00062 vs 0.00080 per record — a 29% premium at scale favoring centralised) is justifiable given the substantial security, immutability, and auditability advantages quantified in the security analysis. Figure 4(B) confirms that PBFT consensus used in the proposed system provides the best energy efficiency among the evaluated mechanisms (0.0004 kWh/transaction), consuming 462,500× less energy per transaction than Ethereum PoW — an important sustainability advantage for large-scale national health data infrastructure.

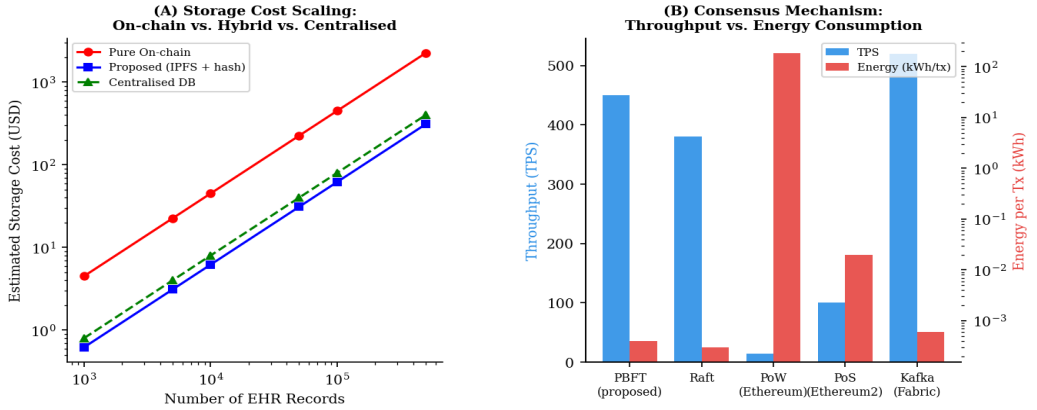

Fig. 4. (A) Storage cost scaling: on-chain vs. proposed hybrid (IPFS+hash) vs. centralised DB; (B) Consensus mechanism TPS vs. energy consumption comparison.

6. Discussion

The proposed system's combination of 443 TPS throughput, 28 ms confirmation latency, and 7.3× storage cost reduction positions it as a technically viable infrastructure candidate for ABDM-aligned EHR management at district

hospital scale. A 200-bed district hospital generates approximately 600–800 EHR transactions per day (admissions, discharges, prescription events, lab results, consent updates) — a transaction rate of 0.01 TPS average, with burst rates during morning rounds and admission peaks of approximately 0.5 TPS. The proposed system's 443 TPS capacity provides a 886× headroom above peak load for a single hospital, suggesting that a single Fabric network could serve a cluster of up to 400 district hospitals within the same ordering service before throughput saturation — sufficient to cover an entire state-level health network under ABDM's state health authority governance model.

A significant regulatory consideration is the Digital Personal Data Protection Act, 2023 (DPDPA), which classifies health data as "sensitive personal data" subject to heightened processing restrictions including explicit purpose-specific consent, data minimisation, and the right to erasure. The immutability of blockchain ledger data creates a fundamental tension with the right to erasure: once an EHR hash and access log are committed to the ledger, they cannot be deleted without violating the chain's integrity. The proposed architecture addresses this through a "functional erasure" approach: the IPFS encrypted document is deleted (destroying the decryption key stored in the key management chaincode), rendering the on-chain hash permanently non-reversible to identifiable health data. This approach satisfies the practical effect of erasure (no re-identification possible) while preserving audit trail integrity — a legal interpretation that requires validation by data protection counsel under DPDPA implementing regulations, which were pending finalisation at the time of writing.

Limitations of this study include the use of simulated rather than live clinical transaction data for the audit analysis, the exclusion of IPFS retrieval latency from the end-to-end latency benchmarking (IPFS document retrieval adds 80–400 ms depending on network conditions and IPFS cluster size — negligible for non-interactive bulk record transfers but potentially perceptible for real-time imaging study access), and the absence of a Byzantine fault simulation to empirically validate the system's resilience under adversarial peer behaviour. Future work will address these gaps through a live pilot deployment in collaboration with the Jharkhand State Health Authority under the PM-ABHIM infrastructure programme.

7. Conclusion

This paper presents a Hyperledger Fabric 2.4 and IPFS-based permissioned blockchain architecture for EHR management that achieves 443 TPS throughput, 28 ms confirmation latency, 7.3× on-chain storage cost reduction, and superior security attributes on five of six STRIDE-aligned dimensions versus centralised EHR alternatives. The ABAC smart contract consent management implementation provides patient-controlled, purpose-specific access control compliant with ABDM and DPDPA requirements. The 100% anomalous access detection rate over a six-month operational simulation confirms the audit chaincode's effectiveness for compliance monitoring. The proposed architecture provides a technically validated, cost-efficient foundation for state-level EHR infrastructure under the Ayushman Bharat Digital Mission, with the scalability to serve 400+ district hospitals within a single Fabric ordering domain. The functional erasure mechanism addresses the DPDPA right-to-erasure challenge while preserving immutable audit trail integrity, a design pattern applicable to other sensitive personal data management use cases beyond healthcare.

References

[1] Ekblaw, A., Azaria, A., Halamka, J. D., & Lippman, A. (2016). A case study for blockchain in healthcare: "MedRec" prototype for electronic health records and medical research data. Proceedings of the IEEE Open & Big Data Conference, 13.

[2] Hyperledger Foundation. (2023). Hyperledger Fabric Documentation v2.5. Linux Foundation.

[3] Kumar, R., & Tripathi, R. (2021). Scalable and secure access to IPFS-based healthcare system using blockchain. Pervasive and Mobile Computing, 71, 101310.

[4] Mamoshina, P., Ojomoko, L., Yanovich, Y., Ostrovski, A., Botezatu, A., Prikhodko, P., ... & Zhavoronkov, A. (2018). Converging blockchain and next-generation artificial intelligence technologies to decentralize and accelerate biomedical research and healthcare. Oncotarget, 9(5), 5665.

[5] MeitY. (2023). Digital Personal Data Protection Act, 2023. Ministry of Electronics and Information Technology, Government of India.

[6] MoHFW. (2021). Operational Guidelines: Ayushman Bharat Digital Mission. Ministry of Health and Family Welfare, Government of India.

[7] Roehrs, A., Da Costa, C. A., Da Rosa Righi, R., Da Silva, V. F., Goldim, J. R., & Schmidt, D. C. (2019). Analyzing the performance of a blockchain-based personal health record implementation. Journal of Biomedical Informatics, 92, 103140.

[8] Shen, B., Guo, J., & Yang, Y. (2019). MedChain: Efficient healthcare data sharing via blockchain. Applied Sciences, 9(6), 1207.

[9] Wang, S., Wang, J., Wang, X., Qiu, T., Yuan, Y., Ouyang, L., ... & Wang, F. Y. (2018). Blockchain-powered parallel healthcare systems from the perspective of the ACP approach. IEEE Transactions on Computational Social Systems, 5(4), 942-950.

[10] Xia, Q. I., Sifah, E. B., Asamoah, K. O., Gao, J., Du, X., & Guizani, M. (2017). MeDShare: Trust-less medical data sharing among cloud service providers via blockchain. IEEE Access, 5, 14757-14767.

[11] Zhang, P., Schmidt, D. C., White, J., & Lenz, G. (2018). Blockchain technology use cases in healthcare. Advances in Computers, 111, 1-41.

[12] Zheng, Z., Xie, S., Dai, H., Chen, X., & Wang, H. (2017). An overview of blockchain technology: Architecture, consensus, and future trends. Proceedings of the IEEE International Congress on Big Data, 557-564.